



Security & Privacy Disclosure, Final Report

Grok for Android (ai.x.grok)

A GDPR Article 22 automated decision-making vector built on users' private X social graph, hardcoded Firebase production credentials with no certificate pinning on an AI platform processing voice, image and conversation data, and seven further undisclosed data flows, unresolved through 91 days and a 28-day-overdue embargo, with zero human reply across six sends and seven automated bot responses

PUBLIC DISCLOSURE, 17 OCTOBER 2026 — EMBARGO REACHED 19 SEPTEMBER 2026, NINE FINDINGS REMAIN UNADDRESSED

Target	xAI Corp (Elon Musk's AI venture) / X Corp, 1355 Market St, San Francisco, CA 94103, USA
Package	ai.x.grok
Platform Audited	Android · v1.1.85-release.01 (versionCode 11185201), 43.7MB base APK, 39,945 smali classes, 7 DEX files, static analysis of the production APK
Initial Disclosure (R1)	2026-06-20
Disclosure / Embargo Date	2026-09-19
Report Published	2026-10-17
Overall Severity	CRITICAL
Findings	9 (2 Critical, 7 High)
Correspondence	6 messages sent by RFI-IRFOS across two email threads to vu.lnerabilities@x.ai, privacy@x.ai, security@x.ai and press@x.ai, 7 automated non-substantive responses received (6 identical privacy-mailbox auto-acknowledgments plus 1 support-ticket reply redirecting to a self-service data-deletion flow, ticket PRIVXAI-7371), 0 human replies ever received
Authorities CC'd	Oesterreichische Datenschutzbehoerde (DSB) · CERT.at · Irish Data Protection Commission

security@rfi-irfos.com • rfi-irfos.com

ZVR 1015608684 • GISA 39261441 • UID ATU83405245 • Steuernummer 68 696/8736 • Graz, Austria

Contents

1	Executive Summary	4
2	Subject & Operator Context	5
3	Findings	5
3.1	C1: Firebase Production Credentials Hardcoded + No Certificate Pinning on an AI Conversation Platform	5
3.2	C2: X Social Graph Personalizes AI Responses — GDPR Article 22 Automated Decision-Making With No Stated Lawful Basis	6
3.3	H1: Intercom Customer Support SDK — 3,505 Classes Routing All Support Data to US Infrastructure	7
3.4	H2: Braze CRM With Location Module — Behavioral Analytics on an AI Chat Platform	7
3.5	H3: AppsFlyer SDK (428 Classes) — Purchase Validation and Attribution on a Paid AI Subscription Platform	7
3.6	H4: GrokVoiceInteractionService — Grok Can Replace the Android System Voice Assistant	8
3.7	H5: SIP/VoIP Stack + MANAGE_OWN_CALLS — Grok Answers Real Incoming Phone Calls	8
3.8	H6: MDM Enterprise Proxy — Corporate IT Can Silently Redirect and Intercept All AI Conversations	8
3.9	H7: Ably Real-Time SDK (448 Classes) + Ashampoo KIM Image Metadata Extraction — Undisclosed Processors	9
4	Impact Analysis	9
5	Response History	9
6	The Complete Chain: Six Messages, Two Threads, Zero Human Replies, 91 Days	10
7	Data Protection — Article by Article	11
8	Regulatory & Legal Landscape	11
9	Indicators of Compromise	12
10	Disclosure Timeline & Outcome	12
11	Residual Risk & Recommendations	13

Executive Summary

Grok for Android (ai.x.grok) is xAI Corp's standalone AI assistant app, offering text and voice chat, image generation and analysis, web search, real phone call answering, and optional replacement of the Android system voice assistant. It is the standalone successor to Grok on X/Twitter and shares xAI's underlying model infrastructure. Static analysis of the production APK (v1.1.85-release.01, 39,945 smali classes) found 2 CRITICAL and 7 HIGH severity findings.

The most significant finding (C2) is a confirmed GDPR Article 22 automated decision-making vector: an explicit configuration string in the production binary confirms that a user's complete X/Twitter social graph, including posts from privately followed protected accounts, can be used to personalize and shape the AI model's responses. No individually stated lawful basis for this personalization exists under Article 22. In parallel (C1), Firebase production credentials are hardcoded in the APK and no network security configuration was found anywhere in the binary, meaning full personal conversation history, voice recordings, uploaded images with embedded location metadata, and payment session data for SuperGrok subscriptions all transit without certificate pinning on any endpoint. Seven further HIGH severity findings document undisclosed data processors (Intercom, Braze with its location module, AppsFlyer, Ably) and undisclosed processing capabilities that go well beyond what a chat app's privacy policy would lead a user to expect: a built-in SIP telephony stack that lets Grok answer real incoming phone calls, a voice-assistant system service that can replace Google Assistant with persistent microphone access, and an enterprise MDM proxy feature that lets a corporate IT administrator silently redirect and intercept all of an employee's AI conversations.

RFI-IRFOS sent the initial disclosure on 2026-06-20 to vulnerabilities@x.ai (cc privacy@x.ai, bcc the Austrian DSB and CERT.at) and followed up four more times over 56 days (2026-06-27, 06-28, 07-23, 08-05, 08-15) to the same and expanded addresses (privacy@x.ai, security@x.ai, press@x.ai), each time receiving the identical automated acknowledgment from xAI's privacy mailbox within seconds. A sixth contact, an automated support-ticket reply (PRIVXAI-7371) received 2026-08-26, redirected the matter to a self-service account-deletion flow and did not engage with a single one of the nine findings or three open questions. At no point across 91 days and seven automated responses did a named human at xAI or X Corp substantively address any finding. The 2026-09-19 embargo, stated explicitly in the original disclosure, has now passed by 28 days and this report is published on the terms stated from the outset.

ID	Severity	Title
C1	CRITICAL	Firebase Production Credentials Hardcoded + No Certificate Pinning on an AI Conversation Platform
C2	CRITICAL	X Social Graph Personalizes AI Responses — GDPR Article 22 Automated Decision-Making With No Stated Lawful Basis
H1	HIGH	Intercom Customer Support SDK — 3,505 Classes Routing All Support Data to US Infrastructure
H2	HIGH	Braze CRM With Location Module — Behavioral Analytics on an AI Chat Platform

ID	Severity	Title
H3	HIGH	AppsFlyer SDK (428 Classes) — Purchase Validation and Attribution on a Paid AI Subscription Platform
H4	HIGH	GrokVoiceInteractionService — Grok Can Replace the Android System Voice Assistant
H5	HIGH	SIP/VoIP Stack + MANAGE_OWN_CALLS — Grok Answers Real Incoming Phone Calls
H6	HIGH	MDM Enterprise Proxy — Corporate IT Can Silently Redirect and Intercept All AI Conversations
H7	HIGH	Ably Real-Time SDK (448 Classes) + Ashampoo KIM Image Metadata Extraction — Undisclosed Processors

Subject & Operator Context

Grok for Android is developed and operated by xAI Corp, with close integration into X Corp's social platform (the app's core personalization feature is built on a user's X/Twitter account data). Both entities share leadership and infrastructure. The app has no disclosed EU establishment; GDPR jurisdiction applies on a per-user, per-member-state basis, with the Austrian DSB and the Irish DPC among the relevant national authorities given X Corp's prior EU operations. The app processes real-time chat, voice, and image data, and offers a paid SuperGrok subscription tier with in-app purchase validation.

Findings

C1: Firebase Production Credentials Hardcoded + No Certificate Pinning on an AI Conversation Platform

CRITICAL (CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:N/VC:H/VI:L/VA:N/SC:N/SI:N/SA:N, 8.7) — hardcoded production credentials combined with a total absence of transport-layer pinning on a platform transmitting conversation, voice and image content

```
google_api_key: AIzaSyA6dlpprRmFg3XPFHs9UXnrUiFcUXX3D6o
google_crash_reporting_api_key: AIzaSyA6dlpprRmFg3XPFHs9UXnrUiFcUXX3D6o
google_app_id: 1:103625933953:android:4e929a8e574ab76308b980
gcm_defaultSenderId: 103625933953
project_id: grok-for-android
google storage bucket: grok-for-android.firebaseioapp
```

```
No network security config.xml found in the production APK.
```

Impact: Firebase credentials (API key, app ID, GCM sender ID, storage bucket, project ID grok-for-android) are hardcoded in res/values/strings.xml. No network_security_config.xml was found anywhere in the production APK. Grok transmits full personal conversation history, voice recordings and voice call sessions, uploaded images and documents including

EXIF/GPS metadata, web search queries, and SuperGrok payment session data. Without certificate pinning on any endpoint, all of these categories are exposed to interception on a compromised or adversarial network. This is the most sensitive data category processed by any app RFI-IRFOS has audited to date, which is why this finding carries a CRITICAL rating despite a CVSS base score reflecting confidentiality impact primarily.

C2: X Social Graph Personalizes AI Responses — GDPR Article 22 Automated Decision-Making With No Stated Lawful Basis

CRITICAL (CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:N/VC:N/VI:N/VA:N/SC:N/SI:H/SA:N, 5.9, blast-radius escalated to CRITICAL) — a user's private social graph, including protected-account posts, shapes general-purpose AI output with no individually disclosed Article 22 basis

```
grok_data_controls_pref_allow_x_personalization_description:
"Allow your X data to be used for personalizing and enhancing your Grok
experience. This data includes: X user profile, X account and location
information, X settings, X preferences, X posts viewable on your X account
(like posts to and from public or protected accounts that you can view),
your conversation history for Grok on X."
```

Impact: RFI-IRFOS blast-radius escalation: raised from the CVSS base rating of MEDIUM/5.9 to CRITICAL because the mechanism documented here is a confirmed GDPR Article 22 fact pattern, automated profiling that shapes a general-purpose AI assistant's output based on a person's private social graph, including posts from protected accounts they privately follow, with no individually stated lawful basis anywhere in the privacy documentation. A CVSS vector built for software vulnerabilities does not capture the societal and individual-rights magnitude of an AI system whose output is systematically filtered by private political, health, and social affiliations the user never consented to have used this way. Users who follow particular political or health-misinformation accounts may receive systematically different AI answers to the same question, and the 'protected accounts' inclusion means private, non-public relationships contribute to that shaping without the account owner's knowledge.

H1: Intercom Customer Support SDK — 3,505 Classes Routing All Support Data to US Infrastructure

HIGH (CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:N/VC:L/VI:N/VA:N/SC:N/SI:N/SA:N, 6.9)

Impact: Intercom (com/intercom/, io/castle/) is the single largest third-party SDK in the APK at 3,505 small classes. All support tickets, in-app chat with Grok support, bug reports, and customer service requests transit through Intercom Inc. (San Francisco, USA), including identity data, device data, and full support ticket content. No transfer mechanism is disclosed at the point of collection, and Intercom is not individually named as a processor anywhere in xAI's privacy documentation.

H2: Braze CRM With Location Module — Behavioral Analytics on an AI Chat Platform

HIGH (CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:N/VC:L/VI:N/VA:N/SC:N/SI:N/SA:N, 6.9)

Impact: Braze (849 classes, com/braze/) ships with its location module (com/braze/location/) alongside push and behavioral-event tracking. On an AI conversation platform this enables geolocation-tagged behavioral profiling of a user's AI usage patterns. braze-html-bridge.js and braze-html-in-app-message-bridge.js confirm Braze's HTML/JS in-app message injection capability. Braze is not individually disclosed as a processor.

H3: AppsFlyer SDK (428 Classes) — Purchase Validation and Attribution on a Paid AI Subscription Platform

HIGH (CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:N/VC:L/VI:N/VA:N/SC:N/SI:N/SA:N, 6.9)

Impact: AppsFlyer (428 classes, com/appsflyer/) ships AFPurchaseDetails, PurchaseHandler and AFInAppEventType, confirming active purchase validation, not merely install attribution, on xAI's SuperGrok subscription billing flow. AppsFlyer Ltd processes device identifiers and purchase event data and is not individually named as a processor or recipient.

H4: GrokVoiceInteractionService — Grok Can Replace the Android System Voice Assistant

HIGH (CVSS v4.0 AV:L/AC:L/AT:N/PR:N/UI:A/VC:H/VI:N/VA:N/SC:N/SI:N/SA:N, 7.1)

Impact: The manifest registers GrokVoiceInteractionService, GrokRecognitionService and GrokVoiceInteractionSessionService, shipped disabled by default but user-activatable to bind Grok as the device's default voice assistant with the same persistent microphone and system-level audio access Google Assistant holds. All subsequent voice queries become xAI processing instead of Google's. The scope of always-on microphone access and the transfer of all voice queries to xAI's US infrastructure in this mode are not individually disclosed.

H5: SIP/VoIP Stack + MANAGE_OWN_CALLS — Grok Answers Real Incoming Phone Calls

HIGH (CVSS v4.0 AV:L/AC:L/AT:N/PR:N/UI:A/VC:H/VI:N/VA:N/SC:N/SI:N/SA:N, 7.1)

Impact: A full SIP telephony stack (android/gov/nist/, android/javax/sip/) ships alongside an android.permission.MANAGE_OWN_CALLS. The app's own strings describe the intended use directly: an AI that picks up calls, talks in a natural voice, and handles questions, bookings and messages on a dedicated number. Telephone call content is among the most sensitive categories of personal communication, and the data processing scope for Grok Voice Agent, what is recorded, stored, and for how long, carries no disclosed Art. 6(1) lawful basis.

H6: MDM Enterprise Proxy — Corporate IT Can Silently Redirect and Intercept All AI Conversations

HIGH (CVSS v4.0 AV:L/AC:L/AT:N/PR:H/UI:N/VC:H/VI:H/VA:N/SC:N/SI:N/SA:N, 7.0)

Impact: An undocumented MDM configuration lets a corporate IT administrator set a single base/proxy URL through which all Grok endpoints, including voice, are routed. No warning is shown to the user when MDM redirection is active. An employee may reasonably believe an AI conversation is private while their employer's infrastructure intercepts it in real time. This capability and its implications for conversation confidentiality are not disclosed anywhere in the privacy documentation.

H7: Ably Real-Time SDK (448 Classes) + Ashampoo KIM Image Metadata Extraction — Undisclosed Processors

HIGH (CVSS v4.0 AV:N/AC:L/AT:N/PR:N/UI:N/VC:L/VI:N/VA:N/SC:N/SI:N/SA:N, 6.9)

Impact: Ably (448 classes, io/ably/) provides the real-time WebSocket/SSE infrastructure streaming all live AI responses through Ably Ltd (UK), undisclosed as a processor. Ashampoo KIM (com/ashampoo/kim/, com/ashampoo/xmp/) reads EXIF and XMP metadata, including GPS coordinates, device model, camera settings and creation timestamp, from every image a user uploads for AI analysis. A user uploading a photo for analysis may have precise location and device history extracted and transmitted alongside the image without this being disclosed anywhere.

Impact Analysis

Grok's central exposure is unusual among the apps in this disclosure series because it is not primarily a data-leakage problem but an output-integrity problem: a general-purpose AI assistant whose answers are shaped by a user's private political and social affiliations, with no disclosure that this is happening, raises questions that go beyond GDPR compliance into the reliability of the tool itself. A user asking Grok a factual question may receive an answer already filtered through accounts they privately follow, without any way to know that filtering occurred. Combined with the SIP call-answering and voice-assistant capabilities, the app's processing footprint extends well past what a 'chat with an AI' description leads a user to expect, into real telephone calls and system-level microphone access. The MDM proxy finding adds an institutional dimension: an employer can redirect and intercept an employee's AI conversations with no on-device warning to either party.

Response History

RFI-IRFOS sent the initial disclosure on 2026-06-20 to vulnerabilities@x.ai (cc privacy@x.ai, bcc the Austrian DSB and CERT.at), stating the findings, three open questions, and the 2026-09-19 disclosure date explicitly. A second send followed 2026-06-27 to the same recipients plus press@x.ai. Three further follow-ups (2026-06-28, 07-23, 08-05, 08-15) went to privacy@x.ai and security@x.ai, each immediately answered by the identical automated acknowledgment from xAI's privacy mailbox: a boilerplate message pointing to the public privacy policy and the data-subject self-service deletion flow, never once addressing a single finding or question. A sixth, later automated reply arrived 2026-08-26 under support ticket PRIVXAI-7371, again redirecting to the self-service deletion flow. Across 91 days, six RFI-IRFOS sends, and seven automated responses, no named human at xAI or X Corp substantively engaged with any of the nine findings.

The Complete Chain: Six Messages, Two Threads, Zero Human Replies, 91 Days

Step 1, disclosure sent: RFI-IRFOS sends the original 2 Critical / 7 High findings to vulnerabilities@x.ai on 20 June 2026, CC privacy@x.ai, BCC the Austrian DSB and CERT.at, stating a 19 September 2026 embargo explicitly. The same-day response is an automated reply from xAI's privacy mailbox, a boilerplate pointer to the public privacy policy and self-service data deletion, unconnected to the disclosure's actual content.

Step 2, broadened recipients: On 27 June 2026, RFI-IRFOS resends to vulnerabilities@x.ai, CC privacy@x.ai and press@x.ai, BCC the Austrian DSB, CERT.at, and the Irish Data Protection Commission. The identical automated reply arrives within seconds.

Step 3, new follow-up thread opens: On 28 June 2026, RFI-IRFOS opens a second thread directly to privacy@x.ai and security@x.ai, BCC the Austrian DSB. Automated reply, identical text.

Step 4, three more follow-ups, same pattern: 23 July 2026 (BCC the Austrian DSB), 5 August 2026 (CC the Austrian DSB), and 15 August 2026, each to privacy@x.ai and security@x.ai, each answered within seconds by the same automated privacy-mailbox acknowledgment, verbatim.

Step 5, a support ticket, still no human: On 26 August 2026, a different automated system responds under support ticket PRIVXAI-7371, this time redirecting to the Grok app's own in-app data-deletion controls and account-deletion flow, a response built for a user who wants their own data erased, not for a formal GDPR security and privacy disclosure naming nine findings.

Step 6, embargo and publication: The stated embargo of 19 September 2026 passes with no human reply ever received. This report publishes 28 days later, on 17 October 2026, on the terms stated from the outset.

Data Protection — Article by Article

Art. 22(1) GDPR, automated decision-making: C2's use of a user's private X social graph, including protected-account posts, to shape AI model output is a confirmed automated-decision-making fact pattern with no individually stated lawful basis, the most significant single finding in this report.

Art. 5(1)(b) GDPR, purpose limitation, and Art. 6(1) GDPR, lawfulness of processing: both apply to C2 in parallel, since data originally shared with Grok on X for one purpose is repurposed to shape a separate standalone assistant's responses without individual consent or disclosure of that repurposing.

Art. 32(1)(a)(b) GDPR, security of processing, and Art. 25(1) GDPR, data protection by design: C1's hardcoded Firebase production credentials and total absence of a network security configuration on a platform carrying conversation, voice and image data are a direct failure of both articles.

Art. 13(1)(e) GDPR, recipients of personal data, and Art. 44-49 GDPR, international transfers: apply to H1 (Intercom, US), H2 (Braze, undisclosed processor and location module), H3 (AppsFlyer, Israel/US), and H7 (Ably, UK, plus Ashampoo KIM's undisclosed EXIF/XMP extraction), none of which are individually named in xAI's privacy documentation.

Art. 5(1)(c) GDPR, data minimisation, and Art. 13(1)(c) GDPR, purposes of processing: apply to H4 (voice assistant microphone access), H5 (SIP call answering and content processing), and H6 (MDM proxy redirection), each a processing capability whose scope is not disclosed to the data subject.

Art. 5(1)(a) GDPR, lawfulness, fairness and transparency: applies additionally to H6, since an employee has no on-device indication that their AI conversations are being routed through an employer-controlled proxy.

Regulatory & Legal Landscape

xAI Corp has no disclosed EU establishment, and X Corp's prior EU presence does not extend a single lead authority to xAI's processing under the one-stop-shop mechanism. GDPR jurisdiction therefore applies on a per-user, per-member-state basis: the Austrian Datenschutzbehörde (DSB) has jurisdiction for Austrian users and was BCC'd from the second message onward, the Irish Data Protection Commission was added from the same point given X Corp's established EU operations, and the UK ICO has jurisdiction for UK users independent of this disclosure's addressing. Art. 22(1) GDPR requires explicit individual disclosure and a stated lawful basis for any automated decision-making that significantly affects a person. Using a user's private Twitter social graph, including protected-account posts, to tune general-purpose AI responses is a novel Art. 22 fact pattern not previously disclosed by xAI in any privacy documentation reviewed during this audit. A Data Protection Impact Assessment under Art. 35 GDPR is warranted for both the X personalization flow (C2) and the voice-call/voice-assistant processing flows (H4, H5), given the scale and sensitivity of the data categories involved.

Indicators of Compromise

The following credentials, service identifiers, and configuration keys were extracted from the production APK (v1.1.85-release.01, versionCode 11185201) via static analysis only. No xAI or X Corp servers, accounts, or infrastructure were probed, and no account was created.

```
C1 -- google_api_key: AIzaSyA6dlpprRmFg3XPFHs9UXnrUiFcUX3D6o
    google_app_id: 1:103625933953:android:4e929a8e574ab76308b980
    project_id: grok-for-android
    No network_security_config.xml present

C2 -- grok_data_controls_pref_allow_x_personalization_description
    (res/values/strings.xml, X social-graph personalization string)

H1 -- com/intercom/, io/castle/, io/castle/highwind (3,505 classes)
H2 -- com/braze/location/, braze-html-bridge.js (849 classes)
H3 -- com/appsflyer/, AFPurchaseDetails, AFInAppEventType (428 classes)
H4 -- ai.x.grok.voice.assistant.system.GrokVoiceInteractionService
    (BIND_VOICE_INTERACTION)
H5 -- android/gov/nist/, android/javax/sip/, MANAGE_OWN_CALLS
H6 -- grok_mdm_base_url_title / grok_mdm_base_url_description
    (corporate proxy base URL configuration)
H7 -- io/ably/ (448 classes), com/ashampoo/kim/, com/ashampoo/xmp/
```

Disclosure Timeline & Outcome

2026-06-20 Initial disclosure sent to vulnerabilities@x.ai, CC privacy@x.ai, BCC the Austrian DSB and CERT.at. Automated privacy-mailbox reply received same day, no substantive engagement.

2026-06-27 Follow-up (7 days) to vulnerabilities@x.ai, CC privacy@x.ai and press@x.ai, BCC the Austrian DSB, CERT.at, and the Irish DPC. Identical automated reply received.

2026-06-28 New follow-up thread opened directly to privacy@x.ai and security@x.ai, BCC the Austrian DSB. Identical automated reply received.

2026-07-23 Follow-up (33 days since original disclosure). Identical automated reply received, no substantive engagement.

2026-08-05 Follow-up (46 days), CC the Austrian DSB. Identical automated reply received.

2026-08-15 Follow-up (56 days), restating all three open questions and naming the response pattern explicitly. Identical automated reply received.

2026-08-26 A separate automated system responds under support ticket PRIVXAI-7371, redirecting to the in-app self-service account-deletion flow. Still no human engagement with any finding.

2026-09-19 Stated embargo date reached. Correspondence record: six RFI-IRFOS messages across two threads, zero bounces, seven automated non-substantive responses, zero human replies.

2026-10-17 This report published, 28 days after the embargo date. All nine findings confirmed unaddressed on the basis of the correspondence record.

Residual Risk & Recommendations

C1 – Immediate: Rotate the hardcoded Firebase production credentials and implement a network security configuration with certificate pinning on every endpoint carrying conversation, voice, image, or payment data.

C2 – Immediate: Document an explicit Art. 22(1) lawful basis for X-social-graph AI response personalization, disclose the mechanism individually in the privacy documentation, and provide users an informed, specific consent control separate from general data-sharing settings.

H1 – Near-term: Name Intercom Inc. individually as a data processor, with its US transfer mechanism, in the privacy documentation.

H2 – Near-term: Disclose the Braze location module as a distinct processing activity with its own stated purpose and legal basis, separate from general push-notification functionality.

H3 – Near-term: Name AppsFlyer Ltd as a processor and recipient for both attribution and purchase-validation data flows.

H4 – Near-term: Disclose the scope of microphone access and voice-query data transfer when Grok is activated as the default system voice assistant, independent of the chat-mode privacy policy.

H5 – Near-term: Document the lawful basis, retention period, and recording/transcription scope for Grok Voice Agent telephone call handling before offering the feature to any user.

H6 – Near-term: Display an on-device, persistent indicator to end users whenever MDM proxy redirection is active, regardless of whether their employer has configured it.

H7 – Near-term: Name Ably Ltd as a processor for real-time response streaming, and document or remove Ashampoo KIM's EXIF/XMP extraction from uploaded images.

Across all findings: establish a monitored intake path for formal security and data-protection disclosures that does not route into the same automated privacy mailbox and self-service deletion flow used for ordinary data-subject requests. Across 91 days, six messages, and seven automated responses, not one of vulnerabilities@x.ai, privacy@x.ai, security@x.ai, or press@x.ai produced a human reply to a disclosure naming a confirmed GDPR Art. 22 automated-decision-making vector, a finding that warrants review independent of the eight further findings in this report.

RFI-IRFOS security@rfi-irfos.com • rfi-irfos.com

ZVR 1015608684 • GISA 39261441 • UID ATU83405245 • Steuernummer 68 696/8736 • Elisabethnergasse 25/10, 8020 Graz, Austria All findings derive from static root-level analysis

of the published application using publicly available tooling. No xAI or X Corp servers, accounts, or infrastructure were probed, and no account was created. Technical evidence is archived under chain of custody and will be provided to regulatory authorities upon lawful request. REF GROK-XAI-2026.